INVESTING IN JAPAN

high Technology

APRIL 1985
$2.00

Vol. 5, No. 4

SPACE STATION
Government and industry launch joint venture

PENTAGON'S INITIATIVE IN ADVANCED COMPUTING

BEYOND BHOPAL: TOWARD "FAIL-SAFE" CHEMICAL PLANTS

James Beggs, Administrator of NASA

0 74820 08405 04

Assessing the STRATEGIC COMPUTING INITIATIVE

DARPA's R&D program pumps millions into industry and academia but raises many concerns

BY DWIGHT B. DAVIS

Ask the average businessperson why the United States is the world leader in computer science, and he or she will likely point to its major university research labs or to an economic system that encourages industrial competition and innovation. But ask a computer scientist the same question, and there is a good chance that the answer will be more specific: the Defense Advanced Research Projects Agency (DARPA). This little publicized agency within the Department of Defense has played a pivotal role in cultivating some of the most important fields in computer science and in pioneering many of the computer-related technologies that now permeate our society. Indeed, the ground-breaking work in such commercially successful applications as timesharing, networking, and computer graphics was performed by university and industry researchers who were funded and directed by DARPA. Thus, although it has received little public attention, DARPA has become well known, respected, and appreciated in the research community.

Slightly over a year ago, DARPA began to assume a higher profile. In October 1983 it unveiled a massive program for dramatic advances in the fields of microelectronics, computer architectures, and artificial intelligence (AI). With projected funding of about $600 million for its first five years, the program—called the Strategic Computing Initiative (SCI)—covers an extremely broad array of technologies that promise to radically alter the role of computers not just in the military but in the commercial world as well. Many computer scientists have warmly embraced the SCI and eagerly anticipate its results—as well as the opportunity to help produce them. Others are more reserved in their support, or even hostile toward the project, because of a number of issues—ethical, economic, and technological—that the SCI evokes.

On its surface, the SCI looks like just one more in a long line of military programs, and a relatively small one at that. Compared with other DOD projects, which can easily run into the billions of dollars, the $150-million-a-year SCI might even seem inconsequential. But in the rarified world of advanced computer science research, the DARPA program is a heavyweight that accounts for more funding than all other sources—government and industry—combined. And because DARPA disburses its funds to university and company researchers, the agency in effect sets the pace and the direction for the development of future computer technologies in the United States. Even IBM, whose annual R&D budget exceeds the revenues of most of its competitors, spends considerably less than DARPA on AI and other highly experimental research, according to Michael Dertouzos, director of MIT's influential Laboratory for Computer Science. (The lab was established as Project MAC in 1963 with DARPA funding.)

Because DARPA's mission is to develop technologies for use by the military, the agency's staff downplays the SCI's likely impact on the commercial sector. But no one disputes that this impact will be real and far-reaching. Already, SCI contracts have been let to firms ranging from Texas Instruments (Dallas) and Martin Marietta (Greenbelt, Md.) to Thinking Machines (Waltham, Mass.), an AI start-up. Historically, DARPA's commercial contractors have moved quickly to incorporate the resultant technologies into products of their own, and there is every reason to suppose that the fruits of the SCI will similarly lend themselves to "dual use" (commercial, as well as military, application). Moreover, future products based on SCI research promise to reshape the computer market in ways that no single company or group of companies could hope to accomplish by itself.

In subject and in scope, the SCI is comparable to Japan's Fifth-Generation Project, which is also pumping millions of dollars into AI and advanced computer research. In fact, many look to the DARPA program as the American computer industry's primary protector against the threat of Japanese domination. Although Robert Cooper, DARPA's director, doesn't promote this viewpoint, he claims that the scientific ambitiousness of the SCI is even greater than that of the Manhattan Project (which developed the first atomic bomb during World War II). In comparison, he says, "the Manhattan Project was a pretty easy job to do." The breadth and difficulty of the SCI have led to some skepticism about DARPA's ability to meet its goals, despite the agency's impressive track record.

The broadest objective of the SCI is to develop machine intelligence to permit the building of both "collaborative" systems, which will closely assist human operators, and autonomous systems, which will function without human intervention. The prospect of autonomous military systems frightens many people, especially those who view the SCI as the first step on the path to computer control of nuclear weapons. People involved with the DARPA project deny that such a state of affairs will ever arise. In fact, many believe that machine intelligence will improve conventional weaponry to the point of reducing the likelihood of nuclear conflict.

But the SCI stirs debate on broader concerns. Many argue that the Department of Defense should not be the most influential player in the development of advanced computer technologies (see "Of swords and plowshares," p. 46). Such influence is seen to be reflected in the three prototype military systems that the SCI has established as goals—an autonomous land vehicle, a pilot's associate, and a battle management system.

Some decry these prototypes, which DARPA maintains will merely demonstrate the utility of the generic technologies they encompass, as a dangerous departure from DARPA's well-proven model for success. "Previously it funded mostly basic research," says Earl Dowdy, a research analyst at the Congressional Office of Technology Assessment. "But the SCI emphasizes applied research and development." And Jim Dray, another OTA research analyst, claims that the SCI "is scaring the wits out of the AI people, because the military goals that were always lurking in the funding have now been made very evident." DARPA personnel, however, vigorously deny that any such shift in emphasis has occurred.

A quiet history of success. DARPA might have become much more visible in its early years (it was founded as ARPA in 1958), but the civilian National Aeronautics and Space Administration (NASA) stole much of the limelight. Both agencies were formed in response to the USSR's successes with Sputnik; it was assumed that NASA would need some time to get up to speed in the space race and that DARPA, by exploiting existing military experience, would fill in the gap. But the civilian space agency developed faster than expected and quickly riveted the public's attention as it moved into manned space flights and moon shots.

DARPA, which is credited with developing technologies that range from phased-array radar to composite materials to laser holography to forward-swept-wing aircraft designs, differs from many other federal research agencies in that it doesn't perform any research itself. Rather, DARPA funds efforts in both academia and industry that are overseen by agency program managers. And even though DARPA is a military entity, its staff emphasizes that the agency's mission is the uncovering and long-term development of new technologies. Once the technologies have reached maturity, DARPA is supposed to back off, leaving it to the patron armed services to decide whether or not to incorporate the technologies into military systems. This way, DARPA has always kept some distance between itself and the military application of its research—distance that has been appreciated by many researchers receiving DARPA funds.

In 1962, four years after its establishment, DARPA acknowledged the importance of the then embryonic field of computer science by forming the Information Processing Techniques Office. Led by a succession of respected office directors, the IPTO funded the development of much of computer science as we know it today. Under its first director, J. C. R. Licklider (now a professor of computer science at MIT), research centered on interactive computing for command and control, with a consequent emphasis on timesharing. DARPA's priorities then shifted to computer net-

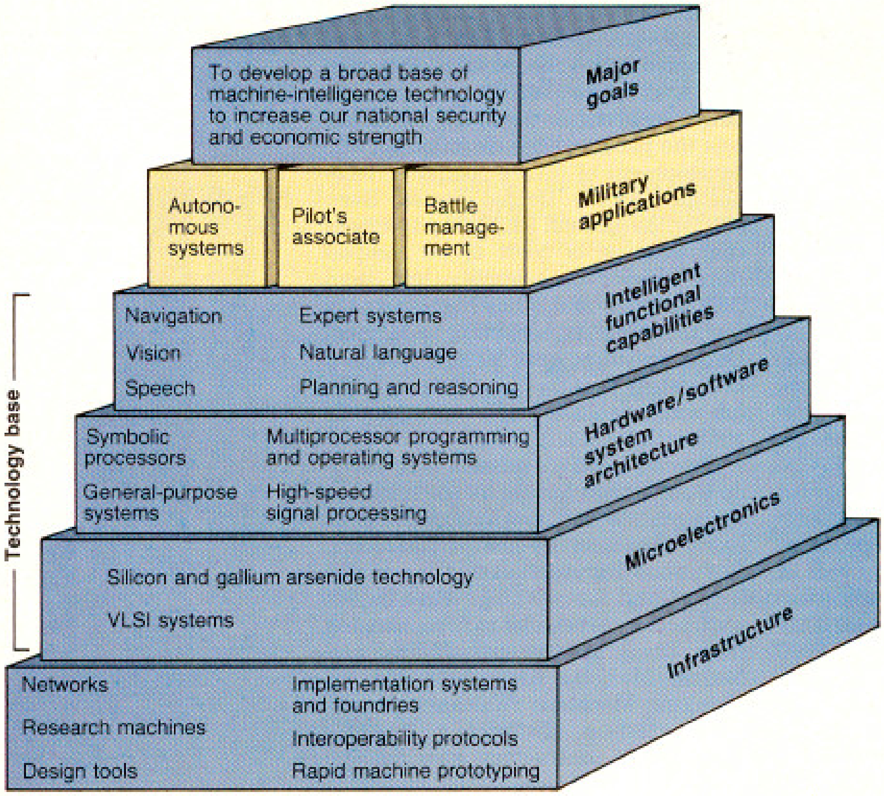

SCI program structure and goals

To achieve its ambitious machine-intelligence goals, DARPA envisions building on an extensive infrastructure, with the advances made at each level serving as a foundation for attaining the goals of the upper levels. Some researchers view the military applications level as a worrisome departure from DARPA's past programs, which usually downplayed or avoided the production of prototype military systems.

DIAGRAMS BY MARK E. ALSOP

works; the most visible development in this area was the ARPANET, a nationwide computer network that has served as the trendsetter for data communications technology and as a research tool in computer science.

The ramifications of these developments have been great, according to MIT's Dertouzos. He maintains that timesharing, networks, and artificial intelligence have made the U.S. a world leader in the computer field. "And DARPA was the key, the only influence in these areas," he says.

Dertouzos grants that two industry-derived developments—microprocessors and personal computers—have also helped shape the United States' leadership. But he insists that DARPA should receive a good deal of the credit for getting the country's computer technology to where it is today.

RICK FRIEDMAN/ BLACK STAR

MIT's Dertouzos: "I could see many of my colleagues, maybe myself 40 years ago, saying we'd refuse to work on radar because it's a weapon of war. Yet today it is also the cornerstone of our world transportation system."

Commercial spinoffs. DARPA's influence has reached beyond the military because of its method of funding. The agency is something of a cross between a corporate research lab and a venture capital operation, says Lynn Conway, chief scientist and assistant director for strategic computing applications at DARPA. In its capacity as a research lab, DARPA attempts to develop technologies of use to the DOD. But rather than doing research inhouse, it funds external projects as if it were a venture capitalist. "It looks for winners, and it bets large sums of money on them," says Conway.

While some of the projects that DARPA funds have no application beyond the military, most involve generic, dual-use technologies. Much of this research is unclassified, and the firms that have developed technologies under contract to DARPA have usually been free to incorporate them into commercial products. Such spinoffs have been so numerous, in fact, that many believe industry has been more adept than the military in exploiting the results of DARPA's research.

No one questions the value of DARPA's commercial spinoffs, but some caution against viewing them as the justification for the agency's existence. Says Jim Lemunyon, staff director of the House Republican Task Force on High Technology Initiatives, "They have a defense mission over there, and they should spend every penny it takes to meet it. But I wouldn't spend any more." Lemunyon believes that nonmilitary agencies, such as the National Science Foundation, should play a heavier role than DARPA in developing basic, dual-use technologies.

Some people believe, however, that DARPA sold its SCI program to Congress as much by stressing its economic benefits to the civilian sector as by detailing its military potential. Sections of the 1983 SCI prospectus, in which the specifics of the initiative are laid out, seem to support this view. Under the heading "Spinoffs from the Technology Base Can Stimulate National Economy," the document states: "The Strategic Computing Program promises the production of machine intelligence technology that will enable yet another major cycle of new economic activity in the computer and electronics industry. . . . Spinoffs from a successful Strategic Computing Program will surge into our industrial community." The document also asserts, "The United States stands to profit greatly both in national security and economic strength by its determination and ability to exploit this new technology."

Nevertheless, DARPA's Cooper denies that the agency has any concerns beyond its military charter. "The spinoffs that occur are not due to any special actions that we take," he says. "They're due only to the fact that the research is generic, unclassified, and available largely to anyone." Still, Cooper admits, commercial spinoff "does happen, we know it happens, and we take advantage of it when we can."

The view that the SCI was formulated as the American response to Japan's Fifth-Generation Project is itself an indication of the program's commercial potential. MIT's Dertouzos notes that he and other scientists such as Stanford's Edward Feigenbaum, alarmed by the Japanese project, expressed their concerns both to industry and to the military establishment. In the civilian sector, such lobbying helped shape the agenda of the industrial consortium called the Microelectronics and Computer Technology Corp. (MCC—Austin, Tex.). But in the federal government, says Dertouzos, the response to the Japanese program eventually coalesced as the SCI. Cooper at DARPA denies this, however; he says it was largely coincidental that the SCI plan followed the announcement of the Fifth-Generation Project. Both programs were bound to be created, he says, because both countries realized that the technologies involved were ripe for intensive development.

With the ultimate goal of bringing work in machine intelligence beyond its current embryonic stage, the SCI will fund research in microelectronic design

DARPA director Cooper: "DARPA can bring sizable resources to bear on a problem, and it can focus and hang in there over a long period of time."

and implementation, parallel processing architectures and software, machine vision, natural-language comprehension, speech recognition, and expert systems. DARPA believes that results in each area will be combined with one another to produce some startling advances.

As evidence of the importance that DARPA places on the development of these and other computer technologies, approximately a quarter of its current $714 million budget is designated for use by the Information Processing Techniques Office. Furthermore, both the overall budget and the proportion allocated to the IPTO are likely to increase. DARPA's budget was recently cut by $300 million when the Strategic Defense Initiative, or "Star Wars," research program was broken off from the agency and established as an independent program. Nevertheless, says Cooper, "I suspect that our budget will get back up to the $1 billion level within the next few years and that the information sciences part of the budget will come up to as much as a third."

The SCI involves such a breadth of technologies that one of the most important aspects of the plan is the development of a massive infrastructure of distributed research centers, supporting agencies, databases, networks, and project managers. Some critics have noted that simply linking various researchers with a network doesn't ensure collaboration or success. But DARPA's reasoning is that "realistic, near-term application demonstrations"—the autonomous land vehicle, the pilot's associate, and the battle management system—will help focus R&D.

OF SWORDS AND PLOWSHARES: *Researchers debate the role of military funding*

The Department of Defense is the United States' single most influential player in setting advanced computer goals and establishing R&D projects. In 1984, for example, almost half of the $104 million in federal funding for basic computer science research came from the DOD, as did about 60% of the $146 million for applied computer research. And if the programmed funding levels for DARPA's Strategic Computing Initiative are realized in subsequent years, the DOD's influence will be greater yet.

Because of the numerous commercial spinoffs that Pentagon funding has engendered over the years, some believe there is no better way to advance computer science than for a DOD agency such as DARPA to administer ambitious programs such as the SCI. Yet many researchers insist that the country would be better served if more computer research funding were available from nonmilitary sources.

For instance, some claim that military funding skews research away from commercial needs. Michael Dertouzos, director of MIT's Laboratory for Computer Science, disputes this charge. History has shown, he says, "that the source of computer science funding and its objectives and aspirations are quite independent of the ultimate payoffs. Science, especially the science yet to be discovered, isn't too happy going where we want to go. It goes where *it* wants to go."

While many agree with the general thrust of Dertouzos's view, few are willing to claim a total independence between research goals and eventual results. "There are many choice points that one passes in traveling through research," notes J. C. R. Licklider, a professor of computer science at MIT and a former DARPA administrator. "And if there's just a little shading and biasing of one to the right or to the left, you wind up in a pretty different place after you go through 1000 more choice points."

Critics also object to military funding of basic research on the grounds that the work is subject to classification and that many skilled scientists become diverted from commercial projects. DARPA discounts both fears. In theory, the bulk of the SCI's basic research will be performed by universities, with most of the applied research done by industry. While the agency admits that some of the applied research might be classified, the generic research—which is of greatest value to academia and industry—will not be. DARPA also admits that the SCI program may put a short-term strain on the availability of trained personnel, but it believes that its funding will actually increase the total base of skilled computer researchers over time.

DARPA's reassurances don't allay everyone's concerns about military funding, but the agency's good reputation helps muffle the opposition. "My political judgment is that the level of U.S. military spending is too high," says John H. Clippinger, president of Brattle Research (Cambridge, Mass.), which is developing products based on AI technology. "But I think that DARPA is the only agency in the government that really funds basic research intelligently."

Ideas abound for ways in which the federal government could improve the effectiveness of its nonmilitary funding for R&D. A presidential commission recently proposed the formation of a cabinet-level Department of Science and Technology that would consolidate the work of several existing agencies. Others call for the formation of a "civilian DARPA" that would follow the successful DARPA model but lack its military charter.

Still others believe that one existing federal agency—the National Science Foundation—should be able to pick up more of the share of basic computer science research, but its track record in this area is not encouraging. The NSF remains far behind DARPA in funding artificial intelligence research and other forefront areas of computer science. MIT's Dertouzos says that the NSF presently devotes only about 2% of its $1.5 billion budget to advanced computer research.

Application goals. The three military prototypes pose staggering technical challenges. Present technology can barely get a robotic vehicle to move slowly from one room to another; according to the Strategic Computing prospectus, the SCI's autonomous land vehicle "must plan routes using digital terrain and environmental data, devise strategies for avoiding unanticipated obstacles, estimate the vehicle's position from landmark and other data, update the on-board digital terrain database, generate moment-to-moment steering and speed commands, and monitor vehicle performance and on-board systems" while traveling at speeds of up to 60 kilometers per hour.

To achieve such performance, the vehicle will need extremely compact, powerful, and rugged computers, which will run a vision system and an expert system for navigation. Almost by definition, these computers will have to be parallel processors, since no computer that executes instructions serially is likely to achieve the necessary performance and compactness. DARPA estimates that the expert system will have to contain as many as 6500 navigation rules, which will fire at a rate of 7000 rules per second. In contrast, current expert systems rarely contain more than 1000 rules or fire at rates faster than 100 rules per second.

The vision system would interpret data from imaging sensors in real time. Aside from requiring much more powerful vision technology than is now available, the system would require a computer that processes from 10 billion to 100 billion instructions per second. Today's supercomputers can attain processing speeds of only a few hundred million instructions per second.

The pilot's associate would be an expert system that assists combat pilots "by off-loading lower-level chores and performing special functions so the pilot may focus his intellectual resources on tactical and strategic objectives." Each pilot could "train" his system to respond in certain ways to his spoken commands, which would be understood by a speech recognition system. The expert system portion of the prototype would have to integrate information from as many as nine "knowledge bases," which would contain facts as well as the rules needed to manipulate the facts. The knowledge bases would cover such areas as tactics and strategy, enemy aircraft, navigation aids, the mission, and enemy defense. The SCI prospectus estimates that monitoring the basic flight systems alone would require several thousand rules.

Meanwhile, the speech recognition system would have to accurately interpret the pilot's spoken commands in the high-noise environment of the cockpit. But current systems can be thrown by problems as simple as slight changes in the speaker's voice because of a cold.

The envisioned battle management system would assist the commander of a naval aircraft carrier during conflicts.

The problem with the NSF, many researchers believe, is that its funding methods are inherently limiting. Whereas DARPA identifies critical research areas, sets goals, hires talented program managers, and gives them a great deal of latitude in funding the best researchers they can find, the NSF mainly responds to unsolicited proposals from the research community. These proposals, moreover, are evaluated by a process of peer review—a mechanism that is well established in the scientific community but that has drawbacks when it comes to producing innovative technologies.

Although peer review promotes honesty, says Dertouzos, "it also promotes the status quo. People continue doing research in the areas in which they're already doing research." Beyond these problems, the NSF simply places more restrictions on its money than does DARPA, he says. "The NSF gives grants of $40,000–50,000, with regulatory restrictions that are devastating. They will fund only 10% of a faculty member's salary, for instance, and they will rarely fund travel. If I had to rely upon NSF money and all its strictures, I would have a dead laboratory."

In the face of these criticisms, there are signs that the NSF is attempting to make some changes. "It is true that the NSF is largely dependent on the proposals sent to it by the science and technology community," admits Roland Schmitt, chairman of the National Science Board, which governs the NSF. "But it can and does try to stimulate proposals in certain areas by launching programs." Schmitt concedes that while the NSF is one of the few federal agencies that haven't adopted a "mission" approach, and is charged with supporting all areas of science, it should probably be devoting a higher percentage of its funds to advanced computer science research.

Even if the NSF shifts its priorities and its methods of funding, or if a "civilian DARPA" is created, some express doubt that any other agency could be as efficient and as productive as DARPA. Craig Fields, director of DARPA's engineering applications office, says he asks critics of the agency the following question: If you had a research

DARPA funding for advanced computer R&D

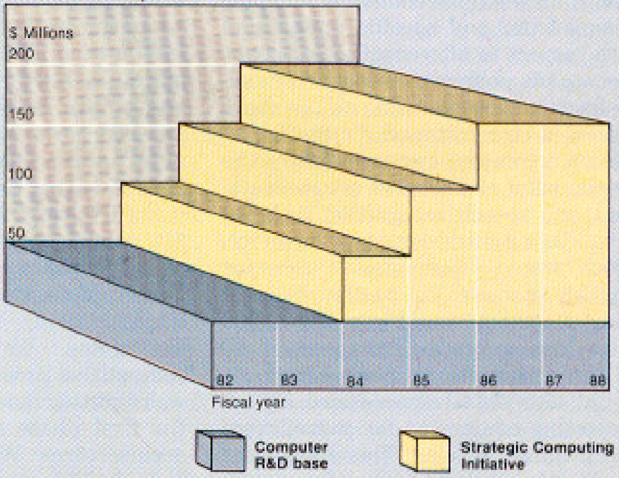

Even before the introduction of the Strategic Computing Initiative, DARPA disbursed about $50 million a year for advanced computer science research. In its first year, the SCI effectively doubled the amount of money DARPA spent in this area. Over its first five years, the initiative is expected to provide $600 million in funding on top of DARPA's existing base.

program optimized for the greatest national economic good, how would it be different in content and organization from the way we're set up? "That's typically the end of the conversation," he says. "We've been a military organization for 26 years, but all the side benefits are evidence to me that the system works."

While he doesn't dispute DARPA's impressive track record, Terry Winograd, associate professor of computer science at Stanford, wonders about the road not taken. "If some agency existed with a different orientation," he says, "what would it have funded instead? It's easy to see the results of what was tried. It's hard to guess what would have happened with something else."

Martin Marietta has received a five-year, $17 million contract to develop advanced computer architectures, artificial intelligence software, and robotic technologies for use in an autonomous land vehicle. An existing vehicle will serve as the platform for the mounting and initial demonstration of these technologies.

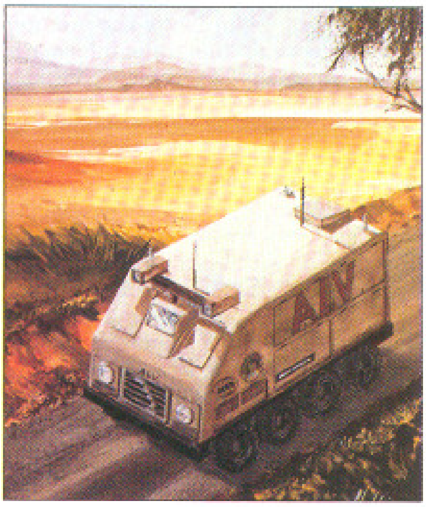

According to the planning document, "The system must alert the decision-maker to the existence of an incipient problem, generate potential responses to the problem in the form of decision options, evaluate these options in the face of uncertainty about the outcome arising from any specific option and with respect to often conflicting goals, execute the preferred option, and monitor its execution, iterating on the above process as circumstances dictate."

Such a complex system requires the development of a number of expert systems and speech recognition systems, as well as a natural-language interface, which would "understand" relevant English statements. DARPA believes such a system would require 20,000 rules and a processing speed of 10 billion instructions per second. It would also entail space-based signal-processing equipment for surveillance and communications. This equipment would use gallium arsenide components because of the technology's low power consumption, high speed, and radiation hardness. But rudimentary gallium arsenide circuits are just beginning to enter the commercial realm.

Widespread skepticism. Not only is the technical scope enormous for each of the SCI's three prototypes, but all are scheduled for completion within 10 years or less. Thus many observers doubt DARPA's ability to make good on its plans. The program's use of detailed timelines, for instance, has led to charges that its success depends upon "scheduled breakthroughs." Some even doubt that certain SCI goals can be achieved in *any* timeframe. At the other extreme, however, are those who perceive no limits whatever on the potential for machine intelligence. "It's impossible to be over-optimistic about what artificial intelligence can achieve," says Marvin Minsky, professor of computer science at MIT and a pioneer in AI. "It's just a question of the timeframe."

For its part, DARPA believes that the SCI is not only feasible but capable of exceeding some of its projections. For example, Craig Fields, director of DARPA's engineering applications office, notes: "We've just finished our competitive procurement process, and we're getting more than we planned" in the first phase of work on computer architectures. Within 30 months, he says, DARPA expects to see the results of nine parallel processing projects, some of which aim to build prototypes that exceed the performance of a Cray supercomputer by two to three orders of magnitude. The architecture program also illustrates a standard SCI strategy: the funding of several simultaneous projects in the same field to engender competition among the various research groups and to serve as a safety mechanism in case some of the individual projects fail.

Still, the skepticism remains. The field of artificial intelligence "has been full of false promises by many of its proponents for many years," says Severo Ornstein, cofounder and chairman of Computer Professionals for Social Responsibility (Palo Alto, Cal.), an organization of 600 or so computer scientists concerned about society's potential misuse of computers. The SCI has become the focus of much of CPSR's concern. "The serious people in the field decry the hype that is being put out right now," says Ornstein. "It's surprising how naive so many computer scientists are. You'd think that after 20 years of failing to do something, they'd understand they're not likely to do it next year."

Global ramifications. What Ornstein finds most worrisome about the SCI, though, is not that its goals may be unrealistic but that it may cause the military to become overdependent on computers, despite their inherent reliability problems. He points to examples of military computers that have given false warnings of missile attacks—warnings that were recognized as false only because of human intervention—and maintains that there is nothing in the SCI prospectus to indicate that the envisioned AI computers will be any more reliable than present machines.

But according to the prospectus, it is precisely the "brittleness" of today's computers that raises the need for the next generation of AI machines. It states that "current computers, having inflexible program logic, are limited in their ability to adapt to unanticipated enemy behavior in the field," and it implies that AI-based computers *could* perform adequately in unpredictable situations.

Critics argue, however, that it is impossible to know how even an intelligent machine will react when faced with a situation that its programmer couldn't anticipate. The contention that AI machines will respond in a "good" way to the unexpected "is a dangerous half-truth," says Terry Winograd, associate professor of computer science at Stanford University. With a rule-based expert system, it is possible that "if some situation comes along that

DARPA's Conway: "Different specialties across computer science and electrical engineering are independently maturing, which creates a startling opportunity to combine and jointly fund things in pursuit of some ambitious goals."

JEANNE MARKLIN / FOLIO

the designer didn't think about, then the rules will combine to do something, as opposed to doing nothing. But if someone is in the middle of a battle and it turns out that the particular combination of events isn't the one that caused the right thing to occur in the test, he doesn't get a second chance."

"Because the system will sometimes do good things that you didn't anticipate, it gives an illusion of flexibility. And that," says Winograd, "is the danger."

DARPA staff members don't dispute this point, but they claim that the flexibility inherent in AI-based computers would allow them to be quickly reprogrammed with new information from the field. "One of the advantages of expert systems is that you can ask them why they arrived at their decisions," Conway explains. "You can chain back through the reasons presented, and if you realize that one of the premises is invalid, you can change it."

But the ability of AI-based computers to adapt on-the-spot to unpredictable situations would be critical for autonomous weapons systems. And if some are merely skeptical about such near-term applications as an autonomous land vehicle, others are terrified by the thought that autonomous systems are being developed essentially to fight a nuclear war. The DARPA staff scoffs at this notion, but others point to a passage in the SCI planning document that they view as ominous: "Commanders remain particularly concerned about the role autonomous systems would play during the transition from peace to hostilities when rules of engagement may be altered quickly. An extremely stressing example of such a case is the projected defense against strategic nuclear missiles, where systems must react so rapidly that it is likely that almost complete reliance will have to be placed on automated systems."

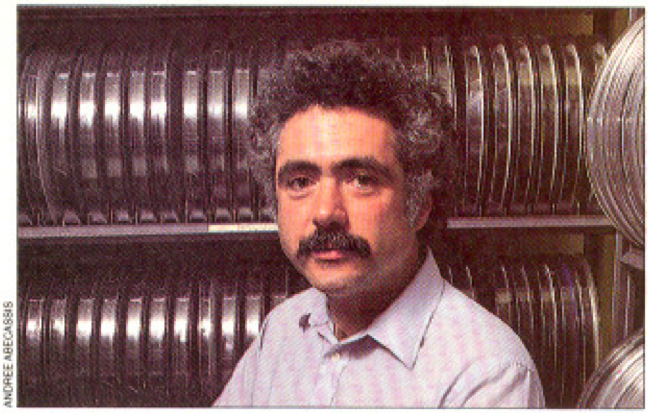
ANDREE ABECASSIS

Stanford's Winograd: "The U.S. is the only Western country that assumes basic research can be done only by the military because nobody else has the power to do it."

On the basis of that passage, CPSR's Ornstein believes that SCI technology would be applicable to the Star Wars missile defense system, which might employ high-energy lasers or particle-beam weapons to shoot down enemy missiles. But DARPA's Cooper replies that because no one yet knows what the proposed Star Wars system would entail, it's impossible to predict the need for autonomous control of the system. Furthermore, he says he is certain that computers will never be placed in a position of controlling nuclear weapons. Last year, Cooper testified before the Senate Foreign Relations Committee about advanced missile surveillance technology. This testimony, he charges, was misinterpreted by many, including the CPSR. "They say I was talking about automatic release of nuclear weapons by machine intelligence. I didn't say that, and I didn't mean that. There will always be several echelons of human judgment between any raw data and the President's decisions."

Stanford's Winograd doesn't question the sincerity of those who, like Cooper, assert that autonomous computers will never control nuclear weapons release. "But that's not DARPA's decision," he says. "They only do research."

Winograd worries that the people who *do* make the decisions about how to apply machine intelligence in military systems are being misled about the technolgy's power. For example, the SCI document implies that AI-based machines already exhibit "common sense"; in fact, a computer with common sense remains something of a Holy Grail in the AI field, and a distant grail at that. "If you tell the generals this thing is intelligent and you don't explain the technology enough," Winograd cautions, "then they will take it at face value and use it in ways that they use people, without understanding the limitations of the machines."

Nevertheless, many in the industrial sector view the project more with enthusiasm than with trepidation. These people, and the companies for which they work, say they are counting on DARPA to continue its tradition of stepping in where other institutions fear to tread. Through the SCI, they believe, DARPA will develop generic technologies for the long-term benefit not only of the military but of technology-based industries and the economy at large. "The countries that are leaders in the information revolution are going to be the dominant economic and geopolitical forces in the world," says MIT's Dertouzos. "And 70–80% of our major discoveries in computer science can be traced to DARPA-supported programs." □

Dwight B. Davis is a senior editor of HIGH TECHNOLOGY.

For further information see **RESOURCES** *on page 84.*

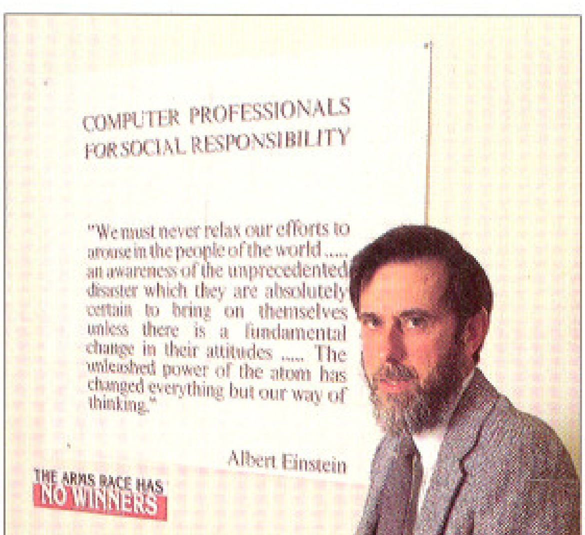

ANDREE ABECASSIS

CPSR chairman Ornstein: "It's clear that defense planners intend to use computers in increasingly sophisticated ways. Ironically, as we rely more and more on them, we become more and more endangered."